The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Uber Drivers *

Cody Cook, Rebecca Diamond, Jonathan Hall

John A. List, and Paul Oyer

December 4, 2017

Abstract

The growth of the "gig" economy generates worker flexibility that, some have speculated, will favor women. We explore gender differences in labor supply choices and earnings in a large gig economy setting: Uber drivers. We first show that this setting is non-discriminatory in that dispatch and pricing decisions are made without conditioning on gender, customer and driver choices (through cancellations of rides) do not affect matches in a gender-biased manner, men and women are equally able when they start driving for Uber, and they learn at an equal rate. Nonetheless, we document a driver hourly earnings gender gap of seven percent which is similar to gaps estimated in other narrowly defined jobs and occupations. We are able to fully explain the gender hourly earnings gap with three factors under the control of each driver: speed, experience, and choice of time and location of work. This suggests that, as the gig economy grows and even in the absence of discrimination, women and men may exercise different choices about the trade-off between earnings and flexibility due to women's relatively high opportunity cost of non-paid-work time and gender-based preference differences.

*We thank seminar participants for valuable input.

Cook and Hall: Uber Technologies, Inc.; Diamond and Oyer: Stanford University Graduate School of Business and NBER; List: University of Chicago, NBER, and Uber Technologies, Inc.

1 Introduction

"The converging roles of men and women are among the grandest advances in society and the economy in the last century. But what must the last chapter contain for there to be equality in the labor market? It must involve changes in the labor market, especially how jobs are structured and remunerated to enhance temporal flexibility. The gender gap in pay would be considerably reduced and might vanish altogether if firms did not have an incentive to disproportionately reward individuals who labored long hours and worked particular hours." Goldin (2014)

The wage gap between men and women has narrowed throughout the past four decades, with 2010 estimates suggesting women earn 88 cents on the dollar as compared to similar men in similar jobs (Blau and Kahn (2017))[1]. Much of the remaining wage gap can be explained by lower hours and weaker continuity of labor force participation by women, especially for middle-age workers where gender wage gaps are largest (Bertrand et al. (2010), Goldin and Katz (2016), Blau and Kahn (2017)). Goldin (2014) has suggested that work hours and disruption in labor force participation dramatically lower wages due to a "job-flexibility" penalty, where imperfect substitution between workers can lead to a convex hours-earnings relationship. In contrast, the role of on-the-job training (Mincer and Polachek (1974)) is thought to play an economically smaller role (Blau and Kahn (2017)).[2]

Recently, many experts have taken an interest in the "gig economy," which can be loosely defined as a collection of labor markets that chunk work into small pieces and then offer those pieces of work to independent workers in real-time with low barriers to entry. Although measuring the economic importance of the gig economy is difficult, estimates suggest about 30% of U.S. workers do some independent work and the share is growing (Oyer (2016) and McKinsey Global Institute (2016)). Female workers may be particularly drawn to gig work because of its flexible hours and transparent compensation.[3]

[1] See Table 4 Panel B of Blau and Kahn (2017), combining the residual wage gap with the effects of experience.

[2] Blau and Kahn (2017) note that the evidence here is mostly based on older studies (Light and Ureta (1995)). Indeed, data on experience often poorly measured contains large amount of measurement error in traditional datasets (Blau and Kahn (2013)).

[3] Hyperwallet (2017) reports that "86% of female gig workers believe gig work offers the opportunity to make equal pay to their male counterparts".

In this paper, we make use of a sample of over a million drivers to quantify the determinants of the gender earnings gap in one of the largest gig economy platforms: Uber's labor spot market. During the period of study, Uber set its fares and fees through a simple, publicly available formula. Importantly, within this setting we know the production and compensation functions, and can therefore identify the specific determinants of driver pay. Further, similar to many parts of the larger gig economy, on Uber there is no negotiation of earnings, earnings are not explicitly based on tenure or hours worked per week, and we can easily rule out customer-side, or dispatch, discrimination. These job attributes explicitly rule out the possibility of a "job-flexibility" wage penalty.[4] We use granular data on Uber drivers and their behaviors in a given hour of the week to precisely measure driver productivity, on-the-job learning, and returns to experience.

We find that men earn roughly 7% more per hour than women on average driving in the same city, which is in line with prior estimates of gender pay gaps within specifically defined jobs (Bayard et al. (2003a), Barth et al. (2017)). Importantly, we find strong evidence that the gap is not due to discrimination on the part of Uber's dispatching software or on the part of riders. We also find that women and men exhibit indistinguishable learnings curves. Rather, we can explain the entire driver gender earnings gap with three factors. First, through the logic of compensating differentials, hourly earnings on Uber vary predictably by location and time of week, and men tend to drive in the more lucrative times (and, to a lesser extent, locations). The second factor is work experience. On-the-job learning is valuable for drivers; we show that drivers on Uber make more as they gain experience in part because they learn where to drive, when to drive, and how to strategically cancel and accept rides. A driver with more than 2,500 lifetime trips completed earns 14% more per hour than a driver who has completed fewer than 100 trips in his or her time on the platform. Even in the relatively simple production of a passenger's ride, there is substantial variation in worker productivity. Male drivers accumulate more experience than women by driving more each week and being less likely to stop driving for Uber. Because of these returns to learning and because the typical male Uber driver has more experience than the typical female - putting them higher on the

[4]This is in contrast with taxi markets in cities such as New York with supply-limiting medallions. In these markets, because the cost of switching drivers is that a valuable medallion will be off the road, contracts are generally structured to make it uneconomical for taxi drivers not to work a very long day. See Haggag et al. (2017)

learning curve - men make more money.

The gender earnings gap that remains after controlling for these two factors can be explained by a single variable: average speed. Increasing speed increases expected driver earnings in almost all Uber settings. Drivers on a trip are paid according to the distance and time they travel on trip and in the vast majority of cases, the extra pay from covering more distance faster outweighs the corresponding loss of time-based pay. We show that men's higher driving speed is due to taste, and not the incentive to drive faster. Drivers of all genders appear insensitive to the incentive to drive faster. Men's higher average speed and the productive value of speed for Uber and the drivers (and, presumably, the passengers) enlarges the gender gap in this labor market.

We interpret these determinants of the gender pay gap - more experience, different times and locations, and higher speed - as taste-based characteristics that are correlated with gender and that make workers more productive. While much prior work has also shown a relationship between the gender gap and factors that are likely to be related to tastes, we know of no prior work that fully explains or fully decomposes the gender pay gap in any setting. Beyond measuring the gender earnings gap and unpacking it completely in an important labor market, our simple analysis provides insights into the roots of the gender pay gap. First, we find that nearly two thirds of the gap can be explained by the return to experience, a factor which is often almost impossible to measure without high frequency data on pay, labor supply, and output. Second, men's willingness to target the most profitable driving hours and locations shows that women continue to pay a cost for fully flexible hours (and locations), even with no convexity in the hours-earning schedule. As the gig economy continues to grow, it will likely bring more flexibility in employment, which is valued by at least some workers (Angrist et al. (2017) and Chen et al. (2017)) document the value of flexibility to drivers) but not by all workers (see Mas and Pallais (2017)). However, the returns to experience and the temporal and geographic variation in worker productivity will likely persist and thus lead to a persistent gender wage gap.

The remainder of the study proceeds as follows. We begin with a brief, and select, summary of the literature studying gender pay gaps. We then describe our data, and show that, using data from nearly two million drivers, there is a roughly 7% gap in hourly earnings between men and

women. Having established that there is a gender earnings gap for drivers, we study the details of how drivers are compensated so that we can break down all components that affect driver pay. We focus on drivers in the Chicago metropolitan area to reveal the primary determinants of the earnings gap.[5] We conclude with implications and summary remarks.

2 Literature

For a general overview of the literature on male/female wage differentials and the factors that lead to them, see Altonji and Blank (1999), Bertrand (2011), and Blau and Kahn (2017). While there are many papers that compare men's and women's earnings across broad groups of the population, a more limited set of papers takes the approach taken here of looking at gender differences within a single employer and/or a narrowly defined set of workers.

Bayard et al. (2003a) uses employer-employee matched data in the US from 1989 to analyze within-establishment gender wage gaps. They find a gender wage gap of 16% within occupations and establishments, which can account for about 50% of the overall wage gap. Since these results are almost three decades old and the economy-wide gender gap has narrowed substantially in the intervening years (Blau and Kahn (2017)), this suggests our Uber gender age gap of 7% is likely not far from the typical within-firm gap.

Several prior papers have shown clear empirical connections between gender pay gaps and factors that are likely related to gender differences in preferences. For example, Bertrand et al. (2010) focus on how pay differences are related to intensity of work. They look at graduates of a single prestigious MBA program. They show that there is a relatively small gender pay gap (about 4%) at graduation which widens considerably with post-MBA experience. The initial gap is driven by women being much less likely to take jobs in finance (which could be due to taste differences or discrimination). The authors show that the increase in the gap can almost be entirely explained by differences in hours worked due to a combination of women working fewer hours per week conditional on working and women being much more likely to have gaps in their careers. There is a (statistically insignificant) unexplained pay gap of 4% after they control for hours worked (at

[5] We have replicated our empirical work in other metropolitan areas and the conclusions are invariably the same.

the time of the observation and over an entire career), industry, MBA grades and classes, and all other observables. Most of the gender gap is explained by hours of work and this can be explained by taste differences related to child rearing.[6]

The mechanism for the hours/earnings connection is unclear as the authors cannot determine whether the female earnings penalty is due to a convex hours-earnings relationship or a learning-by-doing effect. Our results (though in a very different context) are surprisingly similar and our data enable us to quantify the importance of learning-by-doing.

Goldin and Katz (2016) show that hours differences play out very differently in the market for pharmacists. Pharmacists have become increasingly female over time and the gender wage difference is a small 4%. This 4% is completely driven by women who have children, while they find no gender gap within those without children. As compared to the MBAs in Bertrand et al. (2010), the importance of hours and child rearing is economically weaker, although still statistically significant. The institutional setting of pharmacists employment suggest little scope for a "job-flexibility penalty," although empirically it is challenging to distinguish this effect from on-the-job learning and returns to experience.

Azmat and Ferrer (2017) study the gender gap in the market for young lawyers in the United States. They document a large earnings premium for men. Controlling for performance measures such as hours billed and new clients brought in, the gender wage gap becomes statically insignificant (though they cannot reject a gender gap of at least 6%). Gallen (2015) draws similar conclusions from a broad sample of Danish workers, finding that mothers are much less productive than other women or men which explains most of the wage difference they face. Consistent with our conclusions for Uber drivers and those of Bertrand et al. (2010) for MBA's, these papers highlight the importance of gender differences in preferences as driving the gender gap.

Card et al. (2015) relate pay to another gender preference difference — difference in the willingness to bargain. Using matched employee/employer data from Portugal, Card et al. (2015) show that the gender pay gap is exacerbated by proxies for firm-level rents — that is, the gap increases

[6]Similar conclusions can be drawn from the analysis in Barth et al. (2017). They look at the gender gap over careers and by education. They show that the gender gap grows substantially with age for the college educated due to men's pay rising faster within establishments.

when the firm has a high firm-specific pay premium. They also show a dramatic level of sorting by gender into different firms and that the combination of sorting and firm-specific pay levels accounts for 80% of the gender pay gap. They attribute the lower female return to firm-specific pay premia to women bargaining less aggressively. Black and Strahan (2001) show that the gender pay gap declined (and the fraction of women increased) when U.S. banks were deregulated, which they interpret as competitive pressures reducing the banks' surplus available to share with favored (male) employees. Both papers (as well as related results in Hirsch et al. (2010)) show that a large share of the gender pay gap can be attributed to some combination of bargaining and employer discrimination, though it is difficult to clearly distinguish between these channels (as evidenced by the different interpretations of the authors of the two papers).

In addition to Card et al. (2015), other papers look at how the gender pay gap may be related to yet another difference in preferences — different tastes leading to differential sorting into jobs. For example, Gupta and Rothstein (2005) use matched employer/employee data to show that Danish men earned about 34% more per hour than Danish women in 1995. A full set of controls for human capital observables and establishment-by-occupation indicator variables allows them to explain almost 60% of the gap leaving a 14% hourly residual. This pay gap suggests that a large part of gender pay differences can be explained by differential sorting into firms and occupations. However, a large share remains unexplained and the sorting could be (at least partially) driven by discrimination in hiring or elsewhere. Bayard et al. (2003b) reach similar conclusions using U.S. data.

A related area of research includes studies that explore the gender pay gap using experimental techniques. This work varies from examining how gender is related to preferences for incentive schemes in the lab (see, e.g., Gneezy et al. (2003)) and in the field (Flory et al. (2015)), to measuring gender preferences for bargaining and negotiations (see, e.g., Babcock and Laschever (2003), Leibbrandt and List (2015)). The typical data pattern observed is that women disproportionately shy away from competitive work settings, and that men are more likely to negotiate than women. Scholars have used these findings as potential explanations for the portion of the gender pay gap attributed to unobservables in the studies that examine naturally-occurring data.

3 Uber: Background and Data

3.1 The Uber Marketplace

Uber's software connects riders with drivers willing to provide rides at posted prices. Riders can request a trip through a phone app, and this request is then sent to a nearby driver.[7] The driver can either accept or decline the request during a short time window after seeing the rider's location and whether or not a dynamic "surge price multiplier" is in effect. If the driver declines the ride, then the request is sent to another nearby driver. Some products slightly vary this experience. For example, UberPOOL trips may involve picking up multiple riders traveling along a similar route.

Drivers have full discretion regarding when and where they work. Unlike wage and salary workers, drivers do not receive standard employee benefits like overtime or healthcare. A comprehensive discussion of the classification of drivers as independent contractors is out of the scope of this paper, but driver independence is convenient for this study insofar as we do not need to consider differential value for different kinds of compensation beyond monetary compensation and flexibility.

Drivers are paid only when on a trip and according to a fixed formula. For a given trip, the driver earns a base fare plus time and distance rates from pickup to dropoff. In times of imbalanced supply and demand, as manifested by high wait times and few available drivers, a "surge" multiplier greater than one may multiply the time and distance-based fare formula. Importantly, with the exception of some driver promotions which we discuss later, there are generally no explicit returns to tenure (e.g., promotion), convex returns to hours worked (i.e. higher pay for the 50th hour of work in a week than the 20th), or opportunities for earnings discrepancies based on negotiated pay differentials on Uber.

In our analysis, we focus on the UberX and UberPOOL products to ensure that drivers in our data were completing comparable work and faced similar barriers to entry – other Uber products may have alternative pay structures (e.g., UberEATS) or stricter car and license requirements (e.g., UberBLACK).

[7]The dispatch is sent to the nearest driver except occasional situations when there is an opportunity for overall market efficiency gains. For example, if there are two drivers of similar distance, but the closest one is also much closer to a potential incoming dispatch (i.e. another rider looking to request, but not yet requested) the further away driver may receive the dispatch.

3.2 Driver Earnings

For each trip completed, drivers are paid a base fare plus a per-mile and per-minute rate. In Chicago (as of 2017), drivers are paid a $1.70 base fare plus $0.20 per minute and $0.95 per mile for each UberX trip (which are all, at times of high demand, multiplied by a surge multiplier).

Drivers can also earn money from "incentives." For example, drivers may be offered a bounty for completing a set number of trips in a week. Another type of incentive guarantees drivers a certain surge level for trips taken within a given geography and time (e.g. 1.4x all fares in the Chicago Loop during rush hour). While the use of incentives has varied over time, on average they account for under 9% of a driver's earnings.

With all of these components in mind, we formalize the driver's effective hourly pay $p(\cdot)$ for a given trip as

$$p(\cdot) = 60 * \left(\frac{SM \left(r_b + d_1 r_d + 60 * \frac{d_1 r_t}{s} \right) + I}{w + 60 * \frac{d_0 + d_1}{s}} \right) \tag{3.1}$$

where r_b, r_d, and r_t respectively represent the base fare, per-mile, and per-minute rates, SM is the surge multiplier, d_0 is the distance between accepts and pickup, d_1 is the distance on trip, s is speed, w is wait time for dispatch, and I represents the incentive earnings associated with the trip.

For UberPOOL trips – where multiple riders heading in the same direction can ride together – the pay formula treats the chain of trips as a single trip. The driver still receives a base fare for the initial pickup plus a per-mile and per-minute rate.[8] Importantly, pay does not depend on the number of riders in the car. We return to this construction of driver earnings below and isolate how variation in each parameter of Equation 3.1 contributes to gender pay differences for Uber drivers.

3.3 Data

Our data include all driver-weeks for Uber drivers in the U.S. from January 2015 to March 2017. We limit the data to drivers for Uber's "peer-to-peer services," UberX and UberPOOL - drivers

[8] UberPOOL rates are sometimes marginally lower than UberX rates. In Chicago, the per-mile and base fare are identical to UberX, but the per-minute rate is 6 cents lower.

who have completed a trip on other products such as UberXL, UberBLACK, or UberEATS are excluded.[9]

The resulting data include 1,877,252 drivers, 513,417 of whom are female (27.3%). In total, we observe 24.9 million driver-weeks in 196 cities.[10]

For each driver-week, we track total earnings and hours worked. We compute hourly earnings as the total payout in that week divided by hours worked. A driver is considered to be "working" whenever the app is on and available for trips; that is, while on a trip, enroute to a pickup, or available for a dispatch. All earnings are gross earnings. Costs such as gas, car depreciation, and Uber's fees have not been subtracted from the earnings we present.[11] We discuss costs in more depth below.

3.4 Summary Statistics

Table 1 presents summary statistics of driver pay overall by gender. Uber drivers (conditional on being active in the week) gross an average of $375 per week and $21 per hour. More than 60% of those who start driving are no longer active on the platform six months later (though some of these drivers may be on an extended break). Comparing across gender in Table 1, we find a first hint of differences between male and female Uber drivers. Men make nearly 50% more per week than women, which is primarily a reflection of their choice to work 50% more hours per week. On an hourly basis, men make nearly $3/hour more. Part of this, as we will show, is due to composition effects - men work more in cities that pay more. Men are also less likely to leave the platform; as we will discuss further below, this factor strongly contributes to the hourly earnings gap.

[9] UberEats has a different pay structure than ride sharing, paying piece-rate for pickups, dropoffs, and miles driven, and has less stringent vehicle requirements for drivers. Results are consistent with or without UberEats drivers. UberBLACK drivers are commercially licensed and may face large regulatory barriers to entry depending on the city.

[10] We follow Uber's definition of city, which does not always match canonical definitions. For example, the state of New Hampshire is considered a single city.

[11] Uber increased its commission from 20% to 25% in Sept 2015; however, drivers who joined before then were grandfathered in and still pay only 20%. This differentially impacts women, who are more likely to have joined the platform more recently.

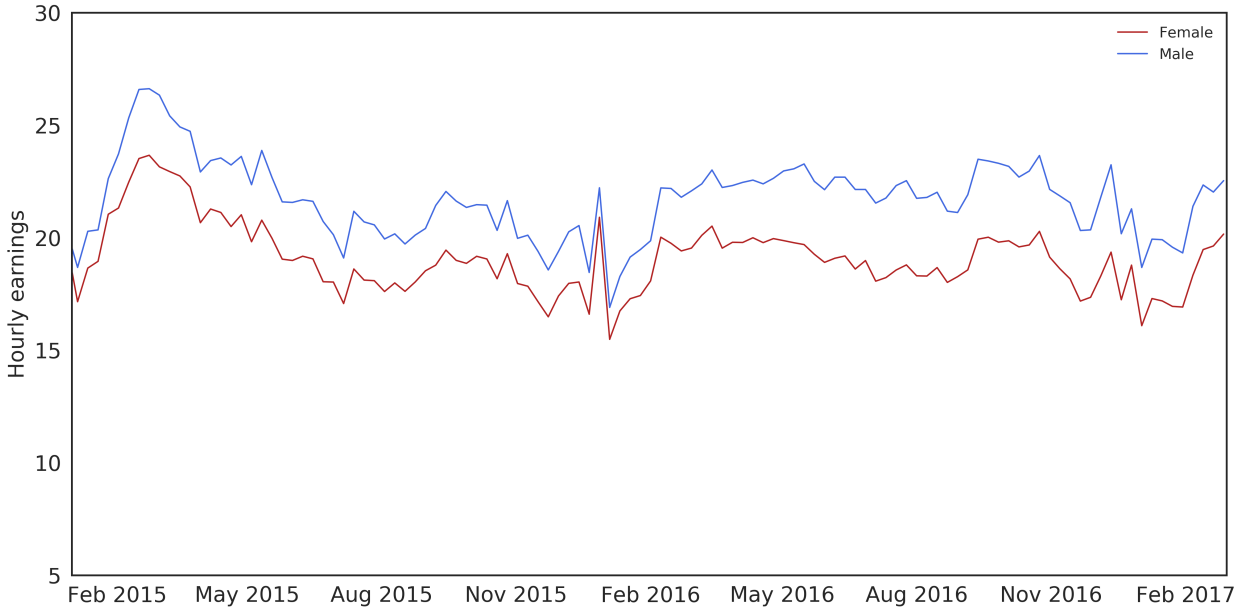
Table 1: Basic summary statistics, all US drivers

	All	Men	Women
Weekly earnings	$375.54	$409.10	$228.75
Hourly earnings	$21.14	$21.66	$18.86
Hours per week	17.02	18.17	12.07
Trips per week	36.15	38.99	25.78
6 month attrition rate	66.6%	63.3%	75.9%
Number of drivers	1,877,252	1,363,835	513,417
Number driver/weeks	24,877,588	20,246,994	4,630,594
Number of Uber trips	740,624,044	646,962,032	93,662,012

Note: Values are based on UberX driver-weeks in the US from January 2015 - March 2017. Average values represent the average across driver-weeks for each week that a driver was active. 6 month attrition rate is the percent of drivers who are no longer active 6 months after the month they completed their first trip

Figure 1 provides a graphical view of the hourly pay gap for all U.S. drivers from early 2015 through early 2017. The $2-3/hour gap seen in Table 1 is fairly constant throughout the sample period. Pay of drivers fluctuates, but the changes are generally gender neutral.

Figure 1: Average hourly earnings, US



Note: Data based on hourly earnings averaged across all driver-weeks in a given week for all UberX and UberPOOL drivers in the US. Earnings are gross; costs such as the Uber commission or gas are not subtracted.

Table 2 uses these national data and measures the Uber driver gender pay gap through a set of standard Mincer regressions. Specifically, we estimate

$$\ln(Earnings_{d,t}) = \beta_0 + \beta_1 isMale_d + \rho \mathbf{X}_{d,t} + \epsilon \tag{3.2}$$

for driver d in time period t, where $Earnings$ are the gross weekly or hourly earnings in that time period, as described above, $isMale$ is an indicator variable for a driver's gender, and $\mathbf{X}$ is a set of controls such as week and city indicator variables.

Table 2: National gender pay gap

	(1) log(weekly earnings)	(2) log(weekly earnings)	(3) log(hourly earnings)	(4) log(hourly earnings)
isMale	0.4142	0.4092	0.0702	0.0653
	(0.002)	(0.002)	(0.001)	(0.001)
Intercept	4.9737	4.9208	2.9280	2.8849
	(0.002)	(0.002)	(0.001)	(0.001)
City	X	X	X	X
Week		X		X
N	24,877,588	24,877,588	24,877,588	24,877,588
Drivers	1,877,252	1,877,252	1,877,252	1,877,252
R^2	0.125	0.136	0.199	0.239

Note: This table documents the gender pay gap for all US cities from January 2015 to March 2017. Data are at the driver-week level; weekly earnings is the entire pay for a given week, while hourly earnings is the pay divided by hours worked in the week. Standard errors (clustered at the driver-level) in parentheses.

Table 2 provides clear evidence that when examining almost two million Uber drivers (representing more than one percent of the United States workforce) across the entire country, but controlling for the city and the conditions for a given week, there remains a large gender pay gap. Men earn about 7% more than women when the analysis is done at the hourly level, indicating that, while a substantial majority of the weekly earnings gender gap is simply due to men driving more hours, there is still a substantial gap when looking at hourly earnings.

This gap may seem surprising: men make 7% more per hour, on average, for doing the same job in a setting where work assignments are made by a gender-blind algorithm and the pay structure is tied directly to output and not negotiated. The 7% differential is as large or larger than hourly differentials in other narrowly defined, relatively homogeneous groups such as recent MBAs (Bertrand

et al. (2010)) and pharmacists (Goldin and Katz (2016)), but are smaller than the differential in economy-wide samples (Blau and Kahn (2017)).

Throughout the rest of the paper, we focus on drivers in Chicago to decompose the gender gap and analyze its economic roots. This choice reduces the dataset to a more tractable size. Table 3 replicates Table 2 but limits the analysis to Chicago drivers. As the table shows, the weekly gender earnings gap in Chicago mirrors the national gap. The hourly Chicago gender earnings gap is somewhat lower, at approximately 5%. This small difference between the national and Chicago gap is due to cross-city differences in the factors that explain the gap. We analyze these factors in detail in Chicago, which provides more insight into the roots of the gap than if we were to focus on the generally small differences across cities. Most importantly, the conclusions we draw are not sensitive to which city we analyze.

Table 3: Chicago gender pay gap

	(1) log(weekly earnings)	(2) log(hourly earnings)
isMale	0.4315	0.0485
	(0.007)	(0.001)
Intercept	5.0487	3.1151
	(0.009)	(0.001)
Week	X	X
N	1,604,627	1,604,627
Drivers	120,019	120,019
R^2	0.038	0.110

Note: This table mirrors Table 2, but limits the data to Chicago drivers from January 2015 to March 2017. Data are at the driver-week level. Hourly earnings are measured as the weekly earnings divided by hours worked in the given week.

4 Decomposing the Wage Gap — Chicago

4.1 Chicago Data and Baseline Gender Gap

By focusing attention on Chicago drivers, we can examine data at the driver-hour, rather than driver-week, level. A driver-hour is defined as a full hour block with some trip activity – for example, 8-9am on a specific Monday. We continue to restrict the data to peer-to-peer drivers in January 2015 to March 2017.

The Chicago dataset includes 120,223 drivers, 36,391 of whom are female (30.2%). In total, we observe 33.0 million driver-hours.[12] As before, we track total gross pay and hours worked for each driver-hour. We compute the implied hourly earnings in a driver-hour as total earnings for trips in that hour divided by minutes worked*60. For trips that span driver-hours, we distribute the pay uniformly between the hours based on the trip time in each hour. In Chicago, certain types of incentive earnings are paid for achieving weekly trips targets, rather than tied to individual trips. We spread these earnings uniformly across minutes worked in the week for which the incentive was earned.

Moving to driver-hour level granularity allows us to control for certain features of a driver's behavior in a given driver-hour. For example, we can now control for where a driver worked, the time of day and day of week, lifetime trips and tenure to-date, and whether the driver rejected a dispatch or canceled a trip that hour.

To control for driving location, we track the "geohash" where a driver is located when he or she accepts a trip. A geohash is a geocoding system that divides the world into a grid of squares of arbitrary precision. For our case, we use geohashes that are approximately three miles by three miles. Within busy areas of Chicago, this is a fairly large area and there may be differences in demand and congestion even within these areas that limit our ability to fully control for geographic effects. We have experimented with finer geographic areas and, given the conclusions do not change, we have not found this worthwhile given the additional computational complexity. To make the computations tractable, we focus on the top fifty (of approximately 570) Chicago geohashes by trip density, which account for 89.2% of trips. The remaining trips are grouped into an "other" bin. For chains of UberPOOL trips, we only include the geohash of the first trip in the chain because a driver does not have control over where to locate for subsequent trips in the chain.

Before using regressions to formally decompose the gender earnings gap shown in Figure 1 (which looks nearly identical when looking only at Chicago), we examine average differences across gender in the factors that determine driver earnings. Recall from Equation 3.1 that driver earnings are a function of wait time between rides, distance to the start of the ride from where the driver

[12]Regressions are run on a 35% subset of drivers. Results are robust to different samples.

accepts it, distance of the ride, speed (both on the ride and on the way to pick up the passenger), the surge rate at the time of the ride, and incentive payments.

Table 4 displays the average of these parameters by gender. Note that these averages are presented on a per-trip basis because that is a more natural way to divide some of the parameters. Table 4 also provides an idea of the sources of the gender pay gap. First, notice that the differences are generally small. Only the difference in incentive payout is more than a few percentage points different and, given the incentive difference is nine cents while the average earnings per trip are about $10, incentives are unlikely to drive the gap. Second, while the individual differences are small, *every* one of the parameters favors men earning more. Men have relatively shorter wait times, shorter trips to the rider, longer rides, faster speed, higher surge, and more incentives.[13] The remainder of our analysis explores which of these differences in Table 4 are important drivers of the Uber gender pay gap and what underlies the differences.

Table 4: Parameter averages

	Men	Women
w – Wait time (min)	10.749	11.022
	(0.005)	(0.011)
d_0 – Accepts-to-pickup distance (mi)	0.485	0.500
	(0.000)	(0.001)
d_1 – Trip distance (mi)	5.035	4.875
	(0.003)	(0.006)
s – Speed (mph)	19.532	18.760
	(0.006)	(0.012)
SM – Surge multiplier	1.051	1.046
	(0.000)	(0.000)
I – Incentive payout ($)	0.903	0.818
	(0.001)	(0.002)
Total per-trip payout ($)	10.142	9.841
	(0.004)	(0.008)

Note: This table documents averages for men and women of the parameters in Equation 3.1. Averages are per-trip based on trips in Chicago between January-February 2017 to avoid issues with seasonality and changes in the composition of driver tenure. Wait time is based on time between either coming online or completing previous trip and picking up passenger for new trip. Trip distance is based on actual route taken; however, accepts-to-pickup distance is the Haversine distance between corresponding coordinates. Standard errors in parentheses.

Table 5 refines the initial Chicago gender earnings gap analysis in Table 3. However, whereas

[13] Equation 3.1 implies that trip distance and speed are ambiguously related to earnings; however, for the values of the other parameters that we observe in the data, earnings are almost always increasing in both distance and speed.

Table 3 utilized weekly observations (the hourly rate in that table is the average hourly rate for a driver in a week) to remain consistent with the regression models using the national data, Table 5 uses driver-hour observations.

Table 5: Baseline gender pay gap

	(1)	(2)	(3)	(4)	(5)	(6)
isMale	0.0356	0.0302	0.0261	0.0220	0.0210	0.0210
	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
riderCancellations						-0.0478
						(0.000)
driverCancellations						-0.0302
						(0.0003)
Intercept	3.0862	3.0912	3.0946	3.0980	3.0989	3.1081
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
Week	X	X	X	X	X	X
Hour of week		X		X	X	X
Geohash			X	X	X	X
Geohash*hour of week					X	X
N	11,572,163	11,572,163	11,572,163	11,572,163	11,572,163	11,572,163
R^2	0.039	0.099	0.092	0.143	0.161	0.164

Note: This table documents the evolution of the gender pay gap as time and location covariates are added. Data are at the driver-hour level. The outcome variable is log of hourly earnings. Hour of week controls for each of 168 hours. Geohash controls are a vector of dummies for whether a driver began a trip in a given geohash. To keep computations tractable, we include the top 50 geohashes in Chicago (which covers almost 90% of trips); trips beginning in other geohashes are captured by an 'other' dummy. Standard errors (clustered at the driver-level) in parentheses.

Column 1 of Table 5 reveals a baseline Chicago gender pay gap of 3.6% at the driver-hour level, controlling only for overall conditions in a given week.[14] Column 2 adds 168 hour-of-week indicator variables. Note that this increases the R-square of the regression quite substantially, which is what we would expect given that Chen et al. (2017) show substantial heterogeneity in Uber driver earnings by hour of week. These hour-of-week controls eliminate 14% of the gender pay gap. This suggests that, while the variation in preferences documented by Chen et al. (2017) are correlated with gender and men drive during relatively high-earnings periods, hour-within-week preference differences are a small part of the gender gap. This suggests that if female Uber drivers receive

[14]This number is lower than the corresponding estimate in Table 3 because the weighting is by driver-hour rather than driver-week. This affects the measured gap for reasons similar to those we discuss below as we decompose the gap.

more non-pecuniary benefits than men from picking which hours to work, they do not pay a large financial price for this flexibility.

Column 3 of Table 5 adds controls for the top fifty Chicago geohashes. This removes about a quarter of the gender gap, indicating that men drive in the parts of Chicago where pay is higher due to factors such as higher surge and shorter waiting times. Combined, the "where and when" variables (geohashes and hour-of-week effects) explain a third of the gender gap (column 4). Column 5 shows that there is no additional explanatory power from the interaction of location and hour, suggesting perhaps surprisingly that the hour-of-week earnings differentials are fairly consistent across areas of Chicago.

While the Uber rider/driver matching algorithm is gender-neutral, *customer* discrimination could contribute to Uber gender pay differences if riders disproportionately cancel rides when paired with a female driver.[15] After requesting a trip, riders see the name and a small image of the driver and can choose to cancel the trip. To determine the potential impact of customer (and driver) discrimination, we control for cancellations in Column 6 of Table 5 and show that customer discrimination does not contribute to the gender pay gap.[16]

Overall, Table 5 shows that time and location explain some of the gender earnings gap but most of it remains unexplained. The remaining gender earnings differential of 2.1% is small compared with overall gender pay gaps measured in the literature, but it is substantial given we are exploring a group of heterogeneous workers doing exactly the same job at the same time and location and being paid by a non-discriminatory algorithm. To identify the roots of the remaining gap, we now turn to gender differences in driver behaviors.

4.2 Returns to experience

As shown by Bertrand et al. (2010), gender differences in hours experience can be an important contributor to gender pay gaps and, as shown by Goldin and Katz (2016), the gender gap can be

[15]Though many studies have hypothesized about customer discrimination and hypothesized that wage residuals may be due to customer preferences (especially race-based discrimination), prior work has not been able to conclusively establish if or when customer discrimination contributes to gender pay gaps.

[16]In the average driver-hour, riders cancel on male drivers 0.1357 times and on female drivers 0.1480 times (8.8% more). However, this is offset by lower cancel rates by women drivers so that the *total* cancellation rate is roughly equal by gender (0.2280 per hour for men and 0.2216 for women)

relatively small if a profession has low penalties for part-time work. Though not examining gender, Haggag et al. (2017) show that learning-by-doing and experience are important for New York City taxi drivers. While Uber drivers may learn in some ways similar to taxi drivers, there are likely important differences. For example, Uber rates fluctuate with surge prices (unlike fixed taxi fares), Uber uses an assignment algorithm to offer rides to drivers, Uber drivers use GPS, and Uber drives are not customarily paid a tip.

We build on this past work on gender differences in experience and on driver learning by measuring how experience affects Uber driver earnings and exploring the relationship between learning, gender, and pay. An important consideration is that while Uber pays according to a fixed formula, many of the parameters of the formula (that is, the variables listed in Table 4) are within the driver's control. For example, drivers can indirectly affect the surge multiplier and wait times by choosing where and when to work and directly affect their driving speed by simply driving faster. As drivers work more, they can begin to learn optimal driving behaviors to maximize earnings.

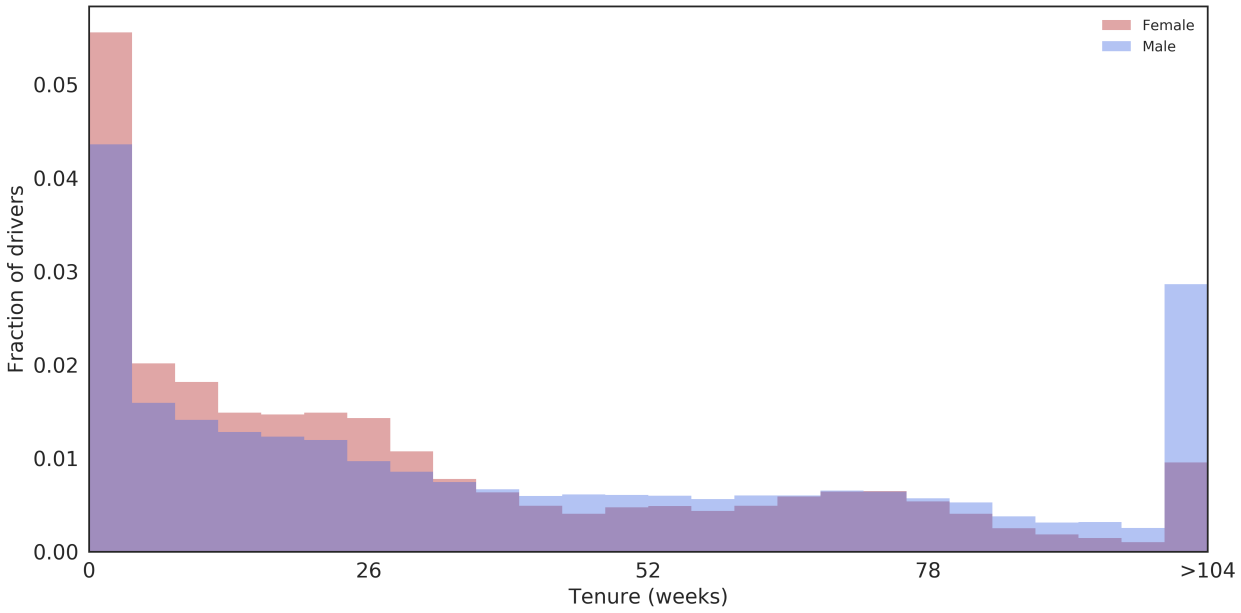
Another activity that may generate a return to experience is "dual apping" - when drivers accept rides from both Uber and a competitor (primarily Lyft). Dual apping has the potential to increase earnings due to less down time and the ability to pick rides at the higher rate if the surge multiplier differs across platforms. Drivers likely become more likely to dual app with experience as they learn from other riders and forums such as Facebook groups. Apps exist that help drivers manage dual apping effectively. We do not have a credible way to determine the degree to which this affects earnings, however.

Figure 2 provides a visual indication of why possible returns to experience can affect the gender earnings gap. The figure, which shows the average tenure of all drivers with a completed trip in January 2017, reveals that men are far more likely to have been driving on Uber for over 2 years. Women are likely to have joined in recent months. Further, Figure 3 shows that men accumulate completed trips at a faster rate than women. Interestingly, just as with Chicago MBAs (Bertrand et al. (2010)), male Uber drivers accumulate more experience by being more likely to work continuously and by working more hours conditional on working.

Figure 4 demonstrates the raw driver returns to experience as measured by cumulative number

of trips driven. There is a clear learning curve, which is especially steep early in a driver's tenure. Drivers continue to learn valuable skills on the job through at least 2,500 trips with a fully experienced driver earning about $3 per hour (more than 10%) more than a driver in his or her first 500 trips. In principle, the rise in earnings shown in Figure 4 could be a selection effect if better drivers last longer on the Uber platform. While there is some degree of selection into staying on the platform based on earnings, Figure 4 looks identical if we limit the graph to drivers that complete at least 1,000 trips or that drive for Uber for at least six months. This suggests the pattern in Figure 4 is a true learning effect.

Figure 2: Distribution of driver tenure, January 2017

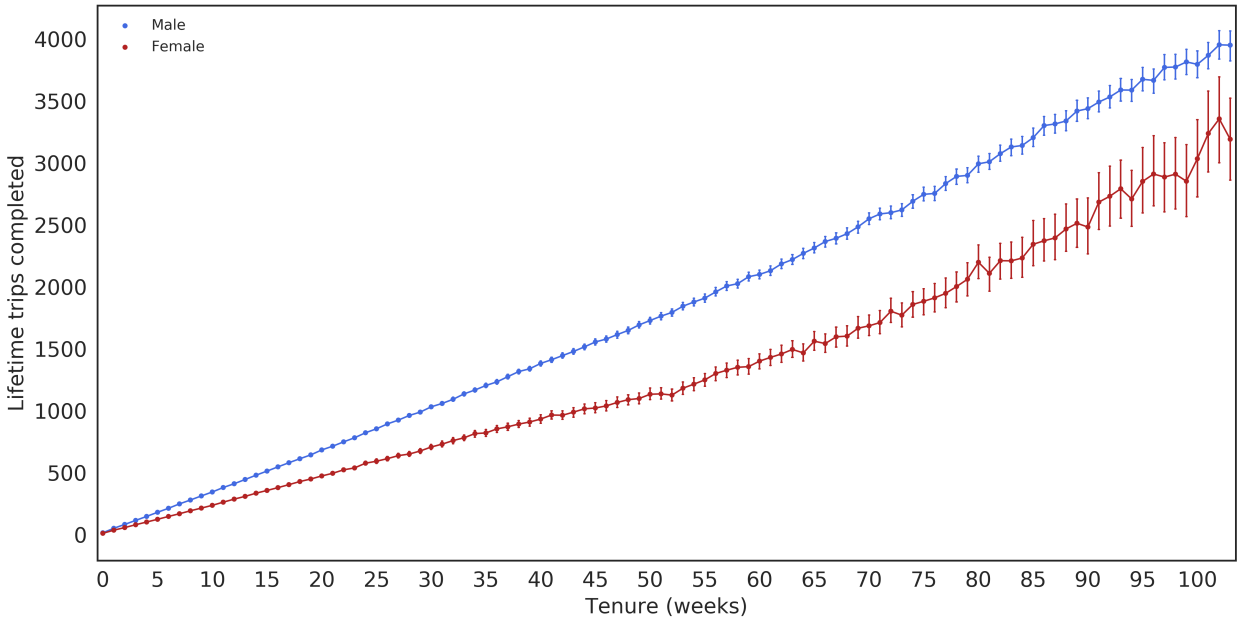


Note: This figure shows the average weeks of tenure for drivers that completed a trip in January 2017; we limit to a single month to avoid composition effects. Tenure is measured as the number of weeks since a driver's first completed trip.

In Table 6, we return to our earnings regression and show that there are substantial returns to experience on Uber. Column 1 shows that drivers who have completed over 2500 trips make nearly 14% more than those in their first 100 trips. Gender imbalances in this learning are clearly important as, controlling for experience, the gender earnings gap shrinks to 1.4% or roughly a third of the initial earnings gap in Chicago.[17]

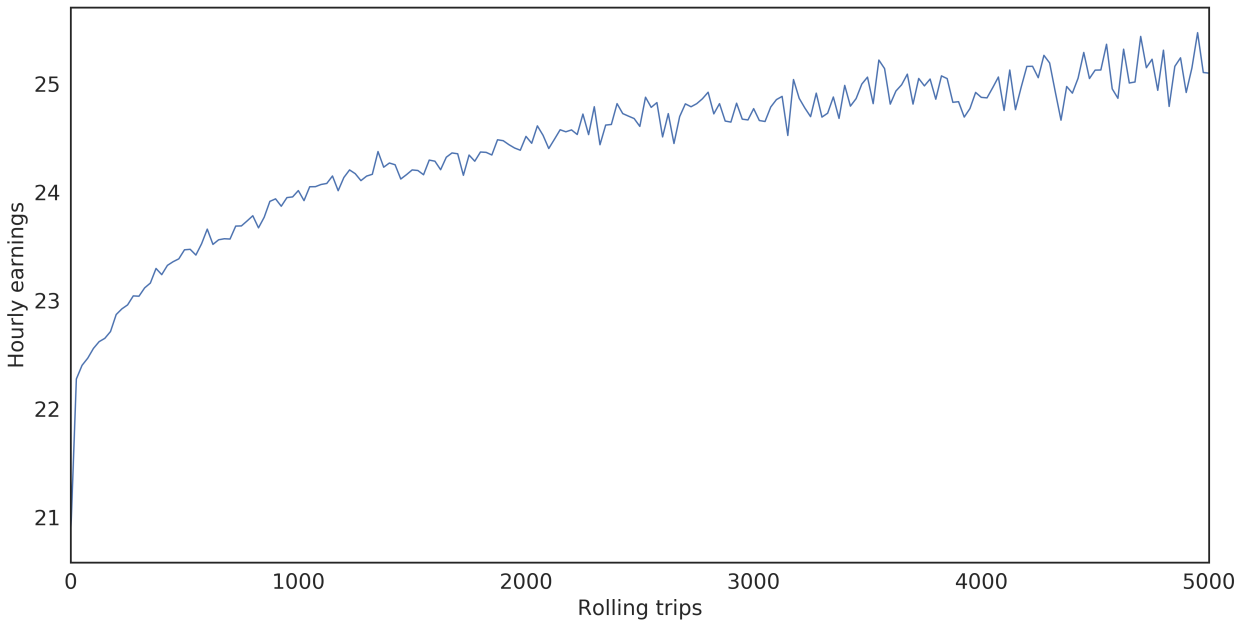
[17]These five categories of experience capture the relevant value of experience. We have experimented with other parametric forms of experience in these regressions and the results are qualitatively similar.

Figure 3: Accumulation of trips over weeks of driving



Note: This figure shows the average number of lifetime trips completed for drivers of a certain tenure. Tenure is based on the number of weeks since a driver completed their first trip. The data only include driver-weeks with >0 trips.

Figure 4: Returns to experience



Note: This figure shows the average earnings of drivers with a given number of rolling trips completed prior to a day of work; rolling trips are binned into buckets of 25 trips completed. Data include all Chicago drivers from January 2015 to March 2017.

With controls for time of week (column 2), the gender gap is further reduced to under 1% but the returns to experience do not change noticably. On the other hand, controls for driver location (column 3) do not reduce the gender gap but substantially reduce the returns to experience. Combined, these two columns suggest that the primary effect of experience on earnings comes from learning *where* to drive and that men and women have differences in terms of their taste for *when* to drive. Overall, the table shows that men and women, through having joined Uber at different times and accumulating experience at different rates, are on different parts of the learning curve and this drives a large part of the gender earnings gap.

In addition to deciding where and when to drive, drivers can affect their earnings through strategic actions. We consider two such strategic actions: rejecting dispatches and canceling rides. When drivers receive a dispatch, they are told where the rider is and the estimated time-to-pickup and can then choose to accept or reject the dispatch. This information can be valuable in assessing the quality of a given dispatch. If a rider is particularly far away, then there is an additional cost; recall that drivers are not compensated for the time it takes to drive to meet a rider.[18] If a driver has reason to think that, by rejecting a ride, he or she will be offered a closer dispatch shortly, that driver may be able to increase expected earnings by not accepting the first dispatch. Savvy drivers will also realize that a high time-to-pickup ride may indicate an imbalance in supply and demand that may soon be corrected by a higher surge.[19]

Once a driver accepts a dispatch, the driver can cancel the trip before picking up the rider. After accepting, drivers are able to contact the rider. Some may do so to learn about the rider destination—for example, calling and asking if the rider is headed to the airport—and canceling if the driver believes the trip will not be worth the time. Experienced drivers may also learn to cancel when they have reason to believe the driver will not show up.

Figure 5 shows the relationship between tenure, acceptance rate of dispatches, and earnings. As inexperienced drivers accept fewer dispatches, their earnings decrease relative to those that accept more. However, as *experienced* drivers accept fewer dispatches, their earnings *increase*. Figure 6

[18] Effective October 2017, Uber initiated a system where drivers are paid (and riders are charged) for particularly long pickups.

[19] Surge rates update every two minutes.

mirrors Figure 5, but for quintiles of cancellation rates. Interestingly, as more experienced drivers cancel more trips, their earnings increase; in contrast, inexperienced drivers' earnings decrease as they cancel a greater number of trips.

Table 7 adds dummy variables that indicate whether a driver rejected a dispatch or canceled a trip during a given driver-hour to our prior regressions. Controlling for time and geography, there is a negative impact on earnings of rejecting a dispatch or canceling a trip. However, this negative effect decreases as experience increases (while still remaining negative). Receiving a bad draw dispatch can never have a positive effect on earnings. A driver either completes the trip, which likely took longer than it was worth, or recognizes that it was a bad draw, rejects or cancels it, and then must wait for the next dispatch. As drivers gain experience, they can more accurately estimate the trade-off between rejecting and having to wait for a new dispatch versus accepting and completing a potentially low value trip.

Even with controls for strategic rejecting and canceling, there are substantial (but smaller) returns to experience. Drivers with over 2500 trips make 6.2% more than those in their first 100 trips. These and earlier regressions show that Uber drivers become more productive (and earn more) as they learn where to drive, when to drive, and how to strategically cancel and accept rides.

While Table 7 sheds light on the driver learning process, these regressions leave the Uber gender pay gap unaffected. There are returns to learning how to reject/cancel and they affect men and women differently because of their relationship to driver experience. But men and women do not differentially accept or reject rides in a way that affects their pay, holding driver experience constant. In all of our analyses, there are no gender differences in the learning process. Learning affects the gender gap because, though each additional ride teaches men and women the same valuable skills, men accumulate ride experience (and, therefore, information) faster than women.

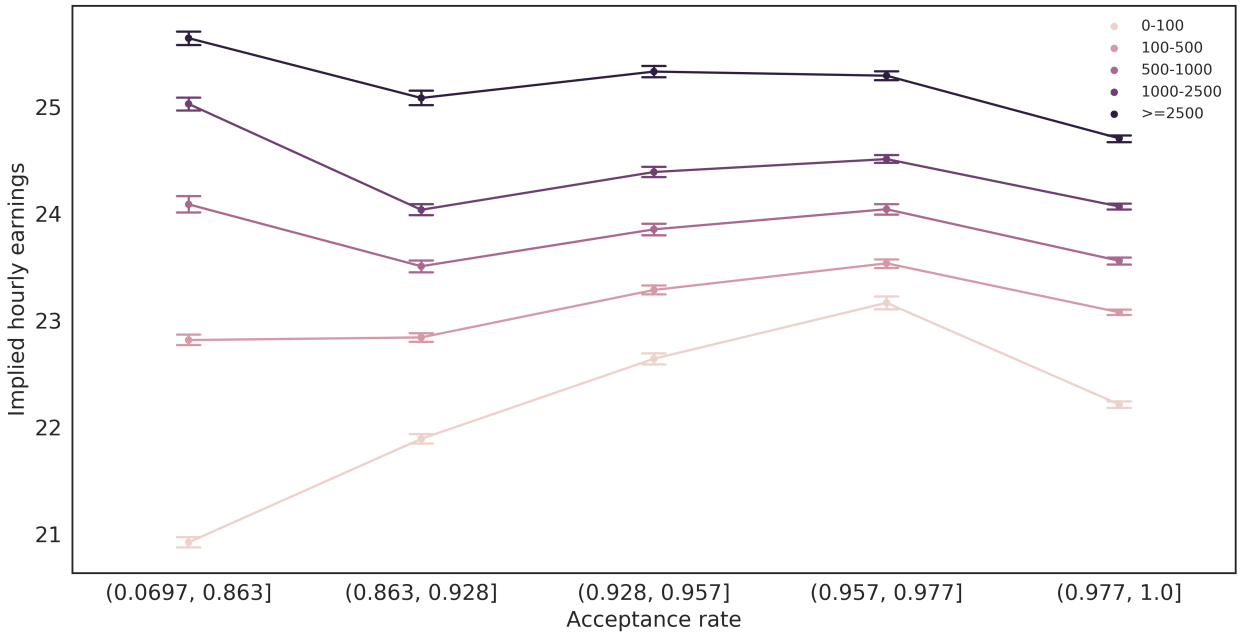
Overall, we conclude that, even in this short-term gig economy environment, experience and gender differences in experience play out in a way that contributes substantially to the gender pay gap as men and women are on different parts of the learning curve. On balance, the relationship between experience and the gender pay gap for drivers is surprisingly similar to at least some traditional job environments.

Table 6: Returns to experience

	(1)	(2)	(3)	(4)
isMale	0.0138	0.0083	0.0129	0.0081
	(0.003)	(0.003)	(0.003)	(0.002)
Trips completed: 100-500	0.0530	0.0497	0.0357	0.0339
	(0.001)	(0.001)	(0.001)	(0.001)
Trips completed: 500-1000	0.0773	0.0747	0.0512	0.0495
	(0.002)	(0.002)	(0.002)	(0.001)
Trips completed: 1000-2500	0.1001	0.0990	0.0650	0.0638
	(0.002)	(0.002)	(0.002)	(0.002)
Trips completed: >2500	0.1391	0.1390	0.0877	0.0860
	(0.004)	(0.003)	(0.003)	(0.003)
Intercept	3.0228	3.0294	3.0528	3.0581
	(0.002)	(0.001)	(0.003)	(0.001)
Week	X	X	X	X
Hour of week		X		X
Geohash			X	X
Geohash*hour of week				X
N	11,572,163	11,572,163	11,572,163	11,572,163
R^2	0.048	0.107	0.096	0.165

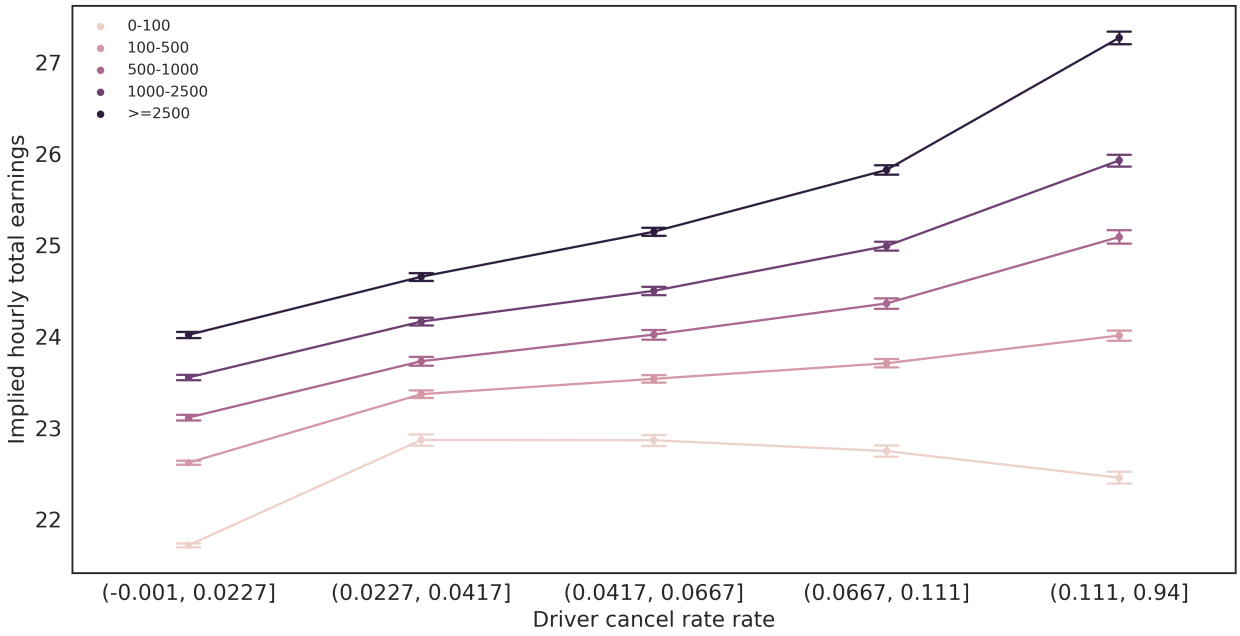
Note: This table expands on the regressions in Table 5 by adding controls for a driver's experience. Experienced is measured as trips completed before a given day of work. Drivers with fewer than 100 completed trips are the excluded category. The outcome variable is log of hourly earnings. Standard errors (clustered at the driver-level) in parentheses.

Figure 5: Returns to experience: rejecting dispatches



Note: This figure documents driver acceptance rate quintiles and corresponding earnings within a given driver-week (for all Chicago driver-weeks with over 5 trips). When a driver receives a dispatch, they are shown pickup location and estimated time to pickup and can choose whether to accept or reject it.

Figure 6: Returns to experience: canceling trips



Note: This figure documents driver cancellation rate quintile and corresponding earnings within a given driver-week (for all Chicago driver-weeks with over 5 trips). A cancellation is when a driver accepts a dispatch and then cancels the trip before picking up the rider. Upon accepting a trip, a driver can contact a rider and inquire about their destination.

Table 7: Returns to strategic rejecting and canceling

	(1)	(2)
isMale	0.0142	0.0093
	(0.003)	(0.002)
Trips completed: 100-500	0.0471	0.0282
	(0.001)	(0.001)
Trips completed: 500-1000	0.0681	0.0414
	(0.002)	(0.001)
Trips completed: 1000-2500	0.0875	0.0542
	(0.002)	(0.002)
Trips completed: >2500	0.1192	0.0619
	(0.004)	(0.003)
rejectDispatch	−0.0757	−0.1151
	(0.001)	(0.001)
rejectDispatch*Trips completed: 100-500	0.0234	0.0211
	(0.002)	(0.002)
rejectDispatch*Trips completed: 500-1000	0.0367	0.0319
	(0.002)	(0.002)
rejectDispatch*Trips completed: 1000-2500	0.0520	0.0430
	(0.003)	(0.002)
rejectDispatch*Trips completed: >2500	0.0765	0.0619
	(0.004)	(0.003)
cancelTrip	−0.0227	−0.0841
	(0.002)	(0.002)
cancelTrip*Trips completed: 100-500	0.0112	0.0131
	(0.003)	(0.003)
cancelTrip*Trips completed: 500-1000	0.0242	0.0275
	(0.004)	(0.003)
cancelTrip*Trips completed: 1000-2500	0.0206	0.0356
	(0.003)	(0.003)
cancelTrip*Trips completed: >2500	0.0462	0.0557
	(0.004)	(0.004)
Intercept	3.0400	3.0872
	(0.003)	(0.002)
Week	X	X
Geohash*hour of week		X
N	11,572,163	11,572,163
R^2	0.049	0.171

Note: This table expands on the regressions in Table 6 by adding covariates for whether a driver rejected a dispatch or canceled a trip in a given hour. Data are at the driver-hour level. The outcome variable is log of hourly earnings. Standard errors (clustered at the driver-level) in parentheses.

4.3 Returns to speed

As shown in Equation 3.1, Uber drivers earn a per-minute and a per-mile rate on each trip. These rates may be balanced such that there is a return to speed. If Uber paid only per-minute, earnings-maximizing drivers would drive as slowly as possible (at least until the disutility of the rider anger exceeds utility from earnings).

The returns to speed will depend on market conditions, such as the expected wait time for a new dispatch. In some unusual circumstances, there are negative returns to speed as the per-minute rates can be relatively valuable if the driver expects to wait a long time until getting another fare. In general, the rates and wait times in our dataset are such that there is a positive expected return to driving faster. This return is fairly steady in our data but it is somewhat higher when driver wait times are shorter. At extremely high speeds, the returns to speed net of costs may turn negative if the risk of a collision or a speeding ticket becomes high enough.

Given that prior research suggests men are more risk tolerant than women and more aggressive (see Bertrand (2011) and, in the context of driving, Dohmen et al. (2011)) and that Table 4 shows that male Uber drivers drive faster than women, we now investigate how driver speed affects the gender earnings gap. Table 8 adds the log of speed as an explanatory variable to our earlier hourly pay regressions. We measure speed as distance on trip divided by time on trip in a given driver-hour.

Control variables are important in this regression because higher pay areas and times of week in Chicago (those areas where there is a more constant stream of fares and where surge is likely to be higher) are also likely more congested, which lowers speed. The coefficient on log speed in Column 1 of Table 8 suggests an elasticity of 27% of speed on earnings — a 1% increase in speed increases earnings by 0.27%. This number increases to 46% in column 2 when we control for geohash (neighborhood) and hour-of-week, thus removing the fact that congestion both lowers speeds and increases earnings. Column 2 shows that controlling for speed and neighborhood reduces our original 3.6% gender gap all the way to just 1%. Adding the learning-by-doing experience variables to this model fully eliminates the gender gap.[20]

[20]The estimate in Column 4 is very precise; the gender pay gap has a 95% confidence interval of -0.6% to 0.2%.

We believe we can describe this speed difference across genders primarily as a difference in taste that happens to have a productive value on Uber rather than a response by male drivers to the incentive to drive faster. First, as aforementioned, others have shown that men are more risk tolerant in general and when driving in particular. Second, when we analyze Uber driver speed as a function of gender, experience, and time/location, we find that men drive 2.2% faster than women.[21]

In addition, we gathered data from the National Highway Travel Survey a nationally-representative survey that gathers demographics, vehicle ownership, and "trip diaries" from 150,000 households. Outside of Uber, there is rarely a pecuniary incentive to drive faster. Despite this, we find that men drive faster in the NHTS sample (details in Appendix Table 17). Gender differences in the taste for speed are a general population phenomenon that have labor market value to Uber drivers.

Table 8: Returns to speed

	(1)	(2)	(3)	(4)
isMale	0.0256	0.0101	0.0016	−0.0019
	(0.004)	(0.002)	(0.002)	(0.002)
logSpeed	0.2677	0.4623	0.2715	0.4616
	(0.002)	(0.001)	(0.002)	(0.001)
Trips completed: 100-500			0.0563	0.0321
			(0.001)	(0.001)
Trips completed: 500-1000			0.0819	0.0460
			(0.002)	(0.001)
Trips completed: 1000-2500			0.1075	0.0594
			(0.003)	(0.002)
Trips completed: >2500			0.1519	0.0810
			(0.004)	(0.003)
Intercept	2.3084	1.7502	2.2293	1.7083
	(0.003)	(0.006)	(0.005)	(0.004)
Week	X	X	X	X
Geohash*hour of week		X		X
N	11,572,163	11,572,163	11,572,163	11,572,163
R^2	0.101	0.282	0.111	0.284

Note: The table expands on earlier regressions by adding log speed as an explanatory variable. Speed is based on total trip distance and duration in a given driver-hour. The outcome variable is log of hourly earnings. Standard errors (clustered at the driver-level) in parentheses.

[21] See Appendix Table 16 for details.

5 Interpretation and Implications

5.1 Summarizing the Decomposition

Using standard pay regressions, we have fully explained the gender gap for Uber drivers. The raw gap in Chicago of approximately four percent can be attributed to three factors: male preference for faster driving, locational choices of drivers, and faster male on-the-job learning. On the one hand, it is somewhat surprising that we can fully explain the gap as we are not aware of prior cross-sectional wage regressions that have precisely and entirely eliminated the gender pay gap in virtually any context. On the other hand, it is also somewhat surprising that there was a gender gap to begin with given Uber rides are allocated algorithmically in a non-discriminatory manner.

To further identify the sources of the differences in pay by gender, we return to our table of averages of all the parameters that affect driver earnings. Table 9 shows the average of each parameter by gender for drivers of three different levels of experience. All the observations are from a short cross-section so as to control for conditions that affect earnings and to eliminate the changes in driver compensation over time.

The table highlights three important themes from our analysis. First, both men and women learn in a productive manner and at roughly the same rate in terms of number of rides. The wait times go down by about 10% over 1,500 rides of experience and are almost identical by gender. Surge rates improve for both genders and are nearly identical for the two genders. Men have slightly longer pickup distances and ride distances throughout, but both genders lower pickup distances and increase trip distances in a similar manner.[22]

Speed is an outlier in that there is not a clear "improvement" over time for drivers. That is because drivers learn that more congested areas are more lucrative. As per our regressions, there is a noteworthy (if not huge) difference in speed by gender that is consistent over tenure.

Table 9 captures the important effects of learning. While men and women learn at the same per-ride *rate*, the driving schedules of men mean that they learn, on average, *more intensively* which generates a gender pay gap. The table also captures that, at all tenures, men prefer to drive faster.

[22]The distance differences seem to be related to men having a stronger preference for airport trips, possibly due to the fact that they work longer shifts and are, therefore, more willing to stray from their base location.

We can more systematically analyze how parameters of the model differ between gender and how that difference feeds into the gender pay gap through Taylor series first order approximations. More precisely, the effect of the difference in parameter v on hourly earnings for men versus women can be approximated by

$$\Delta Earnings_v \approx P_v(\bar{\mathbf{X}})(v_m - v_f) \tag{5.1}$$

where $P_v(\bar{\mathbf{X}})$ is the partial derivative of Equation 3.1 with respect to parameter v evaluated at the average for each parameter, $\bar{\mathbf{X}}$, and $(v_m - v_f)$ is the gender difference in parameter v. We can measure this difference using a regression with the parameter of interest regressed against $isMale$. This allows us to control for experience, where/when, and speed in order to focus on the residual gender-specific gap, i.e. the coefficient on $isMale$, $\hat{\beta_m}$. We run the model both with and without these controls.

Table 10 presents the results of these first order approximations. Due to restrictions on the availability of certain variables [[this sounds sketch...]], these results are based on driver-hours in Chicago between May 2016 to March 2017. Plugging the parameter averages into Equation 3.1 gives us estimated hourly earnings of \$25.261. Comparing this to to the mean hourly earnings in the data of \$25.152 (for the same time frame) confirms that the model accurately captures driver earnings.

Table 10 allows us to examine the levers that men pull in order to earn more money. Without controls, men benefit through each parameter in the earnings function (besides incentives). In order of effect size, men benefit from higher speed (+\$0.403/hr), lower wait times (+\$0.255/hr), higher surge multipliers (+\$0.191/hr), longer trips (+\$0.153/hr), and lower accepts-to-pickup distances (+\$0.007/hr). Men do marginally worse on per-trip incentive earnings (-\$0.004/hr). This nets to an additional \$1.005/hr for men.

As we showed above, when controlling for experience, location/time, and speed (except when endogenous) the earnings gap disappears. However, this is not because men and women now have equal inputs to the earnings function. Men still earn substantially more due to faster speeds (+\$0.198/hr), higher surge (+\$0.082/hr), and lower wait times (+\$0.085/hr); however, that is now negated by men earning less through higher accepts-to-pickup distance (-\$0.158/hr), shorter trips

(-$0.071/hr), and less per-trip incentives (+$0.085). This nets to an additional $0.199 per hour for men; however, the gap is an economically insignificant $0.001 *without* including the returns to speed for men. Since the effects of wait time and surge on pay offset the effects of accepts to pickup distance and trip length, suggests that there exist some gender sorting in when and where to drive within our (hour-of-week) and geohash controls. For example, if there is an event on at a certain time in a specific week, our hour-of-week controls would not account for this. However, this sorting does not contribute to gender differences in wages.

The unique granularity of our data allows us to fully decompose the determinants of gender pay gap. Beyond documenting the factors that explain the gap itself—returns to experience, location and time of week, and speed—we can also examine how these different factors affect inputs into the earnings function of an Uber driver (Equation 3.1). The earnings gap can be fully explained; however, men and women are not completing identical work even once controlling for the factors above. Instead, the differences in inputs to the earnings function cancel each other out; for example, men drive in hours with higher surge, but women earn more in the incentives that Uber offers.

Table 9: Parameter averages by experience

	Men			Women		
Lifetime trips	0-100	700-800	1400-1500	0-100	700-800	1400-1500
w – Wait time (min)	12.086	10.957	10.660	12.052	10.972	10.754
	(0.019)	(0.028)	(0.033)	(0.033)	(0.054)	(0.073)
d_0 – Accepts-to-pickup distance (mi)	0.642	0.535	0.485	0.591	0.506	0.471
	(0.002)	(0.003)	(0.003)	(0.003)	(0.005)	(0.006)
d_1 – Trip distance (mi)	5.068	5.149	5.133	4.811	4.816	4.880
	(0.011)	(0.018)	(0.022)	(0.017)	(0.031)	(0.042)
s – Speed (mph)	20.147	19.798	19.646	19.056	18.551	18.703
	(0.022)	(0.034)	(0.043)	(0.034)	(0.062)	(0.083)
SM – Surge multiplier	1.035	1.046	1.052	1.036	1.044	1.050
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
I – Per trip incentives ($)	0.465	0.847	0.917	0.523	0.789	0.877
	(0.002)	(0.004)	(0.005)	(0.004)	(0.008)	(0.011)

Note: This table documents parameter averages from Equation 3.1 by gender and tenure. Data are limited to trips in Chicago in January-February 2017 to avoid issues with seasonality and changes in composition of driver tenure. Drivers are bucketed based on their lifetime trips before a given day. Wait time is based on time between becoming available for a dispatch (e.g., after coming online or completing a previous trip) and accepting and dispatch. Trip distance is based on actual route taken; however, accepts-to-pickup distance is the Haversine distance between corresponding coordinates.

Table 10: Decomposing the gap: first order approximations

		No controls		Controls	
	Average	$\hat{\beta_m}$	$\Delta Earnings$	$\hat{\beta_m}$	$\Delta Earnings$
w – Wait time (min)	10.189	-0.287	\$0.255	-0.097	\$0.085
d_0 – Accepts-to-pickup (mi)	0.561	-0.003	\$0.007	0.004	\$-0.104
d_1 – Trip distance (mi)	5.238	0.190	\$0.153	-0.070	\$-0.058
s – Speed (mph)	18.860	0.894	\$0.403	0.435	\$0.198
SM – Surge multiplier	1.092	0.009	\$0.191	0.003	\$0.069
I – Incentive payout (\$)	1.127	-0.002	\$-0.004	-0.041	\$-0.085
Week FE		X		X	
Geohash*hour of week				X	
Experience				X	
Speed				X	

Note: This table documents the marginal effect of gender on parameters of Equation 3.1 and how that effect impacts earnings. Data are limited to driver-hours between May 2016 to March 2017 – distance-to-pickup and wait time were not always reliable in earlier data. Coefficients represent the coefficient on $isMale$ from regressions of the form $parameter \sim isMale + controls$. "No controls" includes only controls for the calendar week. "Controls" replicates column (2) in Table 8, with controls for calendar week, experience, location/time, and speed (except when the parameter of interest is speed).

5.2 Costs of driving

We have been using total earnings for Uber drivers as our primary measure of earnings. But, if costs differ in a way that is correlated with gender, we could have understated or overstated the gender "net" pay gap. In some ways, that is not problematic and is consistent with other work. No studies of the gender gap account for differences in costs of working, though the costs of work vary for reasons that may well correlate with gender (such as commuting, clothing, and styling). However, given that a large capital cost is a requirement for independent drivers and that they may deduct some costs from their taxes, it may be appropriate in their case to consider earnings net of direct expenses of driving for hire.

The primary costs Uber drivers face are fuel, maintenance, depreciation, and fines for parking or moving violations. Zoepf et al. (2017) estimate median Uber drivers expenses are 32 cents per mile. They include insurance costs but do not include fines. We will use a conservative estimate of 25 cents per mile for costs other than insurance – Uber covers drivers' insurance costs while driving. A typical Uber driver covers about 20 miles in one hour. The driver earns approximately \$18 net

of Uber's current 25% average share of driver gross earnings. As a "raw" gender gap, we use the 3.56% from column 2 of Table 5 which controls only for calendar week. Based on these numbers, men net 64 cents more per hour than women before expenses. Total average costs per hour – based on 25 cents per mile – are $5, not including fines. Men's costs would have to be ∼13% higher, in terms of fuel, maintenance, depreciation, and fines to erase the gender gap.

First, we consider gasoline costs. Men and women may drive cars with different average fuel efficiency. At a high-level, this appears to be true – of all miles driven by men in the data, 6.4% of them were in a Toyota Prius compared to only 3.6% of miles driven by women. Men have, on average, more incentive to invest in more fuel efficient vehicles due to their longer driving hours. To further test this, we match drivers' vehicles to fuel economy data from the EPA.[23] On average – weighting by miles driven on Uber – women drive cars that get 25.23 miles per gallon in the city and men drive cars that get 26.85 miles per gallon. Men are getting about 6.4% more miles per gallon on average; controlling for gasoline costs would likely increase the gender pay gap.

Another cost to consider is insurance. Though insurance is a large cost for drivers, Uber pays drivers' insurance when they are working. The costs of insurance are relevant, however, as a proxy for accidents (and the downtime that goes with them) and tickets. Men pay more than women for car insurance, though the rates converge at age 26. Accident rates per mile driven are higher for young men than young women but the difference narrows or disappears around age 25. Fatality rates remain higher for men.[24] However, given the insurance rates converge, it seems that the total costs of dangerous driving are about the same by gender after age 25. This suggests that accidents and fines should only vary by gender for drivers under 26. In our sample of Chicago Uber drivers, 15.8% of female drivers and 14.8% of male drivers are under 26. So our gender gap estimates should not be affected by these costs for the vast majority of our sample.

Overall, we cannot estimate differences in costs by gender nearly as precisely as we can estimate gender differences in earnings. However, we also do not see any evidence that the gender pay

[23]Fuel economy data are available at the level of the vehicle make, model, year, and trim. Uber drivers manually enter these fields on sign-up; there are often typos or abbreviations (e.g., "s-class" instead of the exact model). We fuzzy-match based on the Levenshtein distance between the Uber model and the EPA data's model. Results are based only on matches with a Levenstein distance over 0.7 (about 70% of the data). Results are qualitatively similar for different Levenshtein distance thresholds.

[24]See Massie et al. (1995) and Santamarina-Rubioa et al. (2014).

differences are offset by cost differences.

6 Conclusion

The gig economy has become an increasingly large source of earnings for millions of individuals. On Uber alone, over 3.3 million drivers worldwide completed a trip using the platform in August 2017.

Gig economy work is often substantially differentiated from traditional jobs; individuals have more flexibility, are often paid according to a fixed contract, and retain greater control over their earnings. Despite these differences, we show that - much like with traditional jobs - there is a gender pay gap. However, unlike earlier studies, we are able to completely explain the pay gap with three main factors related to driver tastes and learning: returns to experience, a pay premium for faster driving, and preferences for where and when to drive. Indeed, the contribution of the return to experience to gender wage gaps has not gotten much attention in the previous literature, as it is often quite difficult to measure in traditional dataset. We find that even tracking the number of week worked does not accurately quantify experience, as men work more hours per week than women. These results suggest that the role of on-the-job training may contribute to the gender wage gap more broadly in the economy than previously thought.

Uber can, in principle, adjust its driver earnings formula to eliminate the gender pay gap. In order to do that, however, it is necessary to determine which pay gap should be eliminated. Given that speed and experience are productive attributes, one can argue that there is no gender pay gap for Uber drivers holding performance constant. One way to approach the issue is to equalize pay independent of speed or to make earnings equal for people driving at the average speed for each gender, holding other variables constant. A more dramatic approach is to try to equalize earnings by gender without adjusting for the productive effects that come with learning and with driving at times when demand is high, though employers do not generally attempt to do that. Deciding which, if any, of these changes would be best and "fairest" is beyond our goals in this paper.

Overall, our results suggest that, even in the non-discriminatory, transactional, flexible environment of the gig economy, gender-based preferences (especially the value of time not spent at paid

work and, for Uber drivers, a taste for speed) can open gender pay gaps. The preference differences that contribute to pay differences in professional markets for lawyers and MBA's also lead to pay gaps for Uber drivers, suggesting they are pervasive across the skill distribution and whether in the traditional or gig workplace.

References

Altonji, Joseph G. and Rebecca M. Blank, "Race and Gender in the Labor Market," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics, Volume 3*, New York: North-Holland, 1999.

Angrist, Joshua D., Sydnee Caldwell, and Jonathan V. Hall, "Uber vs. Taxi: A DriverâĂŹs Eye View," 2017. NBER Working Paper #23,891.

Azmat, Ghazala and Rosa Ferrer, "Gender Gaps in Performance: Evidence from Young Lawyers," *Journal of Political Economy*, October 2017, *125* (5), 1306–1355.

Babcock, Linda and Sara Laschever, *Women don't ask: Negotiation and the gender divide*, Princeton University Press, 2003.

Barth, Erling, Sari Pekkala Kerr, and Claudia Olivetti, "The Dynamics of Gender Earnings Differentials: Evidence from Establishment Data," 2017. NBER Working Paper 23381.

Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kenneth Troske, "New evidence on sex segregation and sex differences in wages from matched employee-employer data," *Journal of labor Economics*, 2003, *21* (4), 887–922.

_ , _ , _ , and _ , "New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data," *Journal of Labor Economics*, 2003, *21*, 887–922.

Bertrand, Marianne, "New Perspectives on Gender," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics, Volume 4b*, North-Holland, Great Britain, 2011.

_ , Claudia Goldin, and Lawrence F. Katz, "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, July 2010, *2* (3), 228–255.

Black, Sandra E. and Philip E. Strahan, "The Division of Spoils: Rent-Sharing and Discrimination in a Regulated Industry," *American Economic Review*, 2001, *91*, 814–831.

Blau, Francine D and Lawrence M Kahn, "The feasibility and importance of adding measures of actual experience to cross-sectional data collection," *Journal of Labor Economics*, 2013, *31* (S1), S17–S58.

Blau, Francine D. and Lawrence M. Kahn, "The Gender Wage Gap: Extent, Trends, and

Explanations," *Journal of Economic Literature*, 2017, *55*, 789–865.

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *Quarterly Journal of Economics*, 2015, *131*, 633–686.

Chen, M. Keith, Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen, "The Value of Flexible Work: Evidence from Uber Drivers," Working Paper 23296, National Bureau of Economic Research March 2017.

Dohmen, Thomas J., Armin Falk, David Huffman, Juergen Schupp, Uwe Sunde, and Gert G. Wagner, "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences," *Journal of the European Economic Association*, 2011, *9*, 522–550.

Flory, Jeffrey A., Andreas Leibbrandt, and John A. List, "Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions," *The Review of Economic Studies*, 2015, *82* (1), 122–155.

Gallen, Yana, "The Gender Productivity Gap," 2015. Northwestern University.

Gneezy, Uri, Muriel Niederle, and Aldo Rustichini, "Performance in Competitive Environments: Gender Differences*," *The Quarterly Journal of Economics*, 2003, *118* (3), 1049–1074.

Goldin, Claudia, "A grand gender convergence: Its last chapter," *The American Economic Review*, 2014, *104* (4), 1091–1119.

_ **and Lawrence F. Katz**, "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation," *Journal of Labor Economics*, 2016, *34*, 705–746.

Gupta, Nabanita Datta and Donna S. Rothstein, "The Impact of Worker and Establishment-level Characteristics on Male-Female Wage Differentials: Evidence from Danish Matched Employee-Employer Data," *Labour*, 2005, *19*, 1–34.

Haggag, Kareem, Brian McManus, and Giovanni Paci, "Learning by Driving: Productivity Improvements by New York City Taxi Drivers," *American Economic Journal: Applied Economics*, January 2017, *9* (1), 70–95.

Hirsch, Boris, Thorsten Schank, and Claus Schnabel, "Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-

Employee Data from Germany," *Journal of Labor Economics*, 2010, *28*, 291–330.

Hyperwallet, "The Future of Gig Work is Female," 2017. available at www.hyperwallet.com.

Leibbrandt, Andreas and John A. List, "Do Women Avoid Salary Negotiations? Evidence from a Large-Scale Natural Field Experiment," *Management Science*, 2015, *61* (9), 2016–2024.

Light, Audrey and Manuelita Ureta, "Early-career work experience and gender wage differentials," *Journal of labor economics*, 1995, *13* (1), 121–154.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, 2017. forthcoming.

Massie, Dawn L., Kenneth L. Campbell, and Allan F. Williams, "Traffic Accident Involvement Rates by Driver Age and Gender," *Accident Analysis and Prevention*, 1995, *27*, 73–87.

McKinsey Global Institute, "Independent Work: Choice, Necessity, and The Gig Economy," 2016.

Mincer, Jacob and Solomon Polachek, "Family investments in human capital: Earnings of women," *Journal of political Economy*, 1974, *82* (2, Part 2), S76–S108.

Oyer, Paul, "The Independent Workforce in America," 2016. Upwork White Paper.

Santamarina-Rubioa, Elena, Katherine PÃłreza, Marta Olabarriaa, and Ana M. Novoaa, "Gender differences in road traffic injury rate using time travelled as a measure of exposure," *Accident Analysis and Prevention*, 2014, *65*, 1–7.

Zoepf, Stephen, Stella Chen, Paa Adu, and Gonzalo Pozo, "The Economics of Ride Hailing: Driver Expenses, Income and Taxes," 2017. Stanford University.

7 Appendix

7.1 Results in other cities

We repeat our analyses for drivers between January 2015 and March 2017 in Boston, Detroit, and Houston. For the duration of our sample, Uber's main ridesharing competitor Lyft did not operate in Houston, making the city of particular interest.

Tables 11 through 13 present results for four specifications in each city: (1) baseline, (2) adding controls for location and time of week, (3) adding controls for experience, and (4) adding controls for speed. The results for each city tell a similar story: there is a small 3-5% baseline gender pay gap in each city, which can be explained by differences in where/when drivers work, different levels of experience, and preferences for driving speed. In the case of Houston, the pay gap actually reverses once controlling for those three factors and men make an estimated 1.4% *less* per hour than women.

Table 11: Gender pay gap: Boston

	(1)	(2)	(3)	(4)
isMale	0.0493	0.0305	0.0090	−0.0020
	(0.004)	(0.003)	(0.003)	(0.003)
Trips completed: 100-500			0.0742	0.0805
			(0.002)	(0.002)
Trips completed: 500-1000			0.0933	0.0978
			(0.002)	(0.002)
Trips completed: 1000-2500			0.1107	0.1135
			(0.002)	(0.002)
Trips completed: >2500			0.1361	0.1378
			(0.003)	(0.003)
logSpeed				0.4303
				(0.001)
Intercept	3.2016	3.2198	3.1432	1.9150
	(0.004)	(0.003)	(0.003)	(0.005)
Week	X	X	X	X
Geohash*hour of week		X	X	X
N	13,912,058	13,912,058	13,912,058	13,912,058
R^2	0.048	0.200	0.207	0.305

Note: Data include all driver-hours for a random sample of 40% of UberX/Pool drivers in Boston, January 2015 to March 2017. Data are sampled to make the computations more tractable. Geohash*hour of week indicates controls for geohash, hour of week, and the interaction between the two. Tenure is measured as lifetime trips completed before a given day of work. Speed is measured as trip distance divided by trip duration in a given driver-hour. Standard errors (clustered at the driver-level) in parentheses.

7.2 Differential returns to learning

[Men and women don't learn at different rates]

Table 12: Gender pay gap: Detroit

	(1)	(2)	(3)	(4)
isMale	0.0361	0.0164	0.0102	0.0010
	(0.003)	(0.002)	(0.003)	(0.003)
Trips completed: 100-500			0.0323	0.0196
			(0.002)	(0.002)
Trips completed: 500-1000			0.0471	0.0300
			(0.003)	(0.003)
Trips completed: 1000-2500			0.0576	0.0379
			(0.004)	(0.003)
Trips completed: >2500			0.0719	0.0499
			(0.006)	(0.006)
logSpeed				0.6082
				(0.001)
Intercept	2.8938	2.9106	2.8833	0.8608
	(0.004)	(0.003)	(0.002)	(0.005)
Week	X	X	X	X
Geohash*hour of week		X	X	X
N	3,807,809	3,807,809	3,807,809	3,807,809
R^2	0.038	0.142	0.143	0.344

Note: Data include all driver-hours for UberX/Pool drivers in Houston, January 2015 to March 2017. Geohash*hour of week indicates controls for geohash, hour of week, and the interaction between the two. Tenure is measured as lifetime trips completed before a given day of work. Speed is measured as trip distance divided by trip duration in a given driver-hour. Standard errors (clustered at the driver-level) in parentheses.

Table 13: Gender pay gap: Houston

	(1)	(2)	(3)	(4)
isMale	0.0327	0.0154	0.0027	−0.0142
	(0.004)	(0.003)	(0.003)	(0.002)
Trips completed: 100-500			0.0450	0.0237
			(0.001)	(0.001)
Trips completed: 500-1000			0.0675	0.0302
			(0.002)	(0.002)
Trips completed: 1000-2500			0.0834	0.0344
			(0.002)	(0.002)
Trips completed: >2500			0.0919	0.0288
			(0.003)	(0.003)
logSpeed				0.7492
				(0.001)
Intercept	2.8728	2.8877	2.8441	0.4107
	(0.004)	(0.003)	(0.003)	(0.005)
Week	X	X	X	X
Geohash*hour of week		X	X	X
N	6,770,183	6,770,183	6,770,183	6,770,183
R^2	0.021	0.112	0.114	0.368

Note: Data include all driver-hours for all UberX/Pool drivers in Detroit, January 2015 to March 2017. Geohash*hour of week indicates controls for geohash, hour of week, and the interaction between the two. Tenure is measured as lifetime trips completed before a given day of work. Speed is measured as trip distance divided by trip duration in a given driver-hour. Standard errors (clustered at the driver-level) in parentheses.

Table 14: Differential learning

	(1)	(2)
isMale	0.0145	0.0096
	(0.002)	(0.002)
Trips completed: 100-500	0.0529	0.0343
	(0.003)	(0.002)
Trips completed: 500-1000	0.0768	0.0493
	(0.004)	(0.003)
Trips completed: 1000-2500	0.0995	0.0655
	(0.006)	(0.004)
Trips completed: >2500	0.1453	0.0919
	(0.014)	(0.009)
isMale*Trips completed: 100-500	−0.0004	−0.0006
	(0.003)	(0.003)
isMale*Trips completed: 500-1000	0.0004	−0.0002
	(0.004)	(0.004)
isMale*Trips completed: 1000-2500	0.0006	−0.0022
	(0.006)	(0.006)
isMale*Trips completed: >2500	−0.0069	−0.0067
	(0.0012)	(0.012)
Intercept	3.0223	3.0571
	(0.002)	(0.002)
Week	X	X
Geohash*hour of week		X
N	11,572,163	11,572,163
R^2	0.047	0.164

Note: The table expands on results presented in Table 6 by adding interacting gender and experience. Outcome variable is log of hourly earnings. Standard errors (clustered at the driver-level) in parentheses.

7.3 Returns to experience

The returns to experience we document could be driven by selection bias – for example, perhaps drivers with lower earnings drop off the platform and those that reach the higher tenure bins had always been earning more per hour. To test this, we add driver fixed effects to our model. Results are presented in table 15. Adding driver fixed effects attenuates the learning curve by about ∼30% with time and location controls and ∼35% with only controls for the calendar week.

There are other forms of selection bias that could affect how we measure the returns to experience. For example, perhaps drivers construct an expectation of future earnings growth and drop off the platform if the do not expect to continue learning.

Table 15: Returns to experience, driver fixed effects

	(1)	(2)	(3)	(4)
Trips completed: 100-500	0.0539	0.0349	0.0344	0.0242
	(0.001)	(0.001)	(0.001)	(0.001)
Trips completed: 500-1000	0.0788	0.0527	0.0503	0.0374
	(0.002)	(0.002)	(0.002)	(0.002)
Trips completed: 1000-2500	0.1023	0.0688	0.0650	0.0495
	(0.002)	(0.002)	(0.002)	(0.002)
Trips completed: >2500	0.1421	0.0925	0.0876	0.0668
	(0.004)	(0.003)	(0.003)	(0.003)
Intercept	3.0326	3.0614	3.0640	3.0771
	(0.001)	(0.002)	(0.001)	(0.002)
Week	X	X	X	X
Geohash*hour of week				X
Driver		X		X
N	11,572,163	11,572,163	11,572,163	11,572,163
R^2	0.048	0.123	0.165	0.205

Note: This table includes our baseline returns to experience specifications with and without driver fixed effects. Outcome variable is log of hourly earnings. Standard errors (clustered at the driver-level) in parentheses.

7.4 Driving speed

We model driving speed against gender to test whether men drive faster after controlling for experience, location, and hour of week. Results presented in Table 16 show that men drive 2.2% faster

after controls. Table 17 presents similar results based on data from the National Household Travel Survey; even in contexts where there is no pecuniary incentive to drive faster, men still do so.

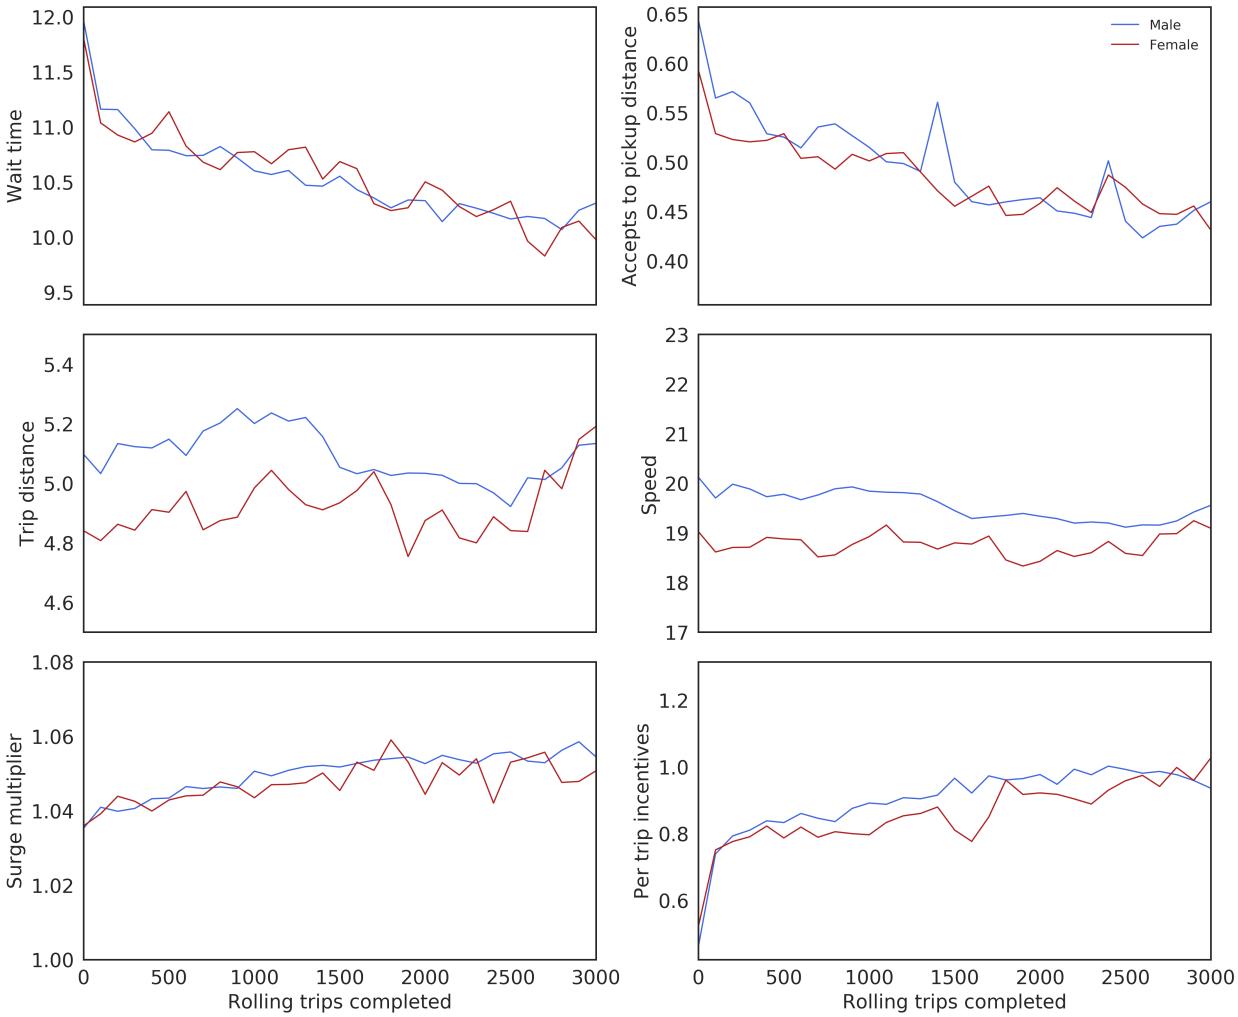
Table 16: Effect of gender on driving speed

	(1)	(2)
isMale	0.0236	0.0218
	(0.002)	(0.002)
Trips completed: 100-500		0.0039
		(0.001)
Trips completed: 500-1000		0.0075
		(0.001)
Trips completed: 1000-2500		0.0096
		(0.002)
Trips completed: >2500		0.0110
		(0.002)
Intercept	2.9174	2.9119
	(0.001)	(0.001)
Week	X	X
Geohash*hour of week	X	X
N	11,572,163	11,572,163
R^2	0.352	0.352

Note: This table regresses log speed in a given driver-hour against the driver's gender and experience. Speed is measured as distance traveled on-trip in an hour over duration on-trip. Standard errors (clustered at the driver-level) in parentheses.

7.5 Additional graphs & tables

Figure 7: Parameter averages by gender and experience



Note: This figure documents parameter averages from Equation 3.1 by gender and tenure. Data are limited to trips in Chicago in January-February 2017 to avoid issues with seasonality and changes in composition of driver tenure. Drivers are bucketed based on their lifetime trips before a given day, binned into 100s. Wait time is based on time between becoming available for a dispatch (e.g., after coming online or completing a previous trip) and accepting and dispatch. Trip distance is based on actual route taken; however, accepts-to-pickup distance is the Haversine distance between corresponding coordinates.

Table 17: Effect of gender on driving speed: NHTS data

	(1)	(2)	(3)	(4)
Male	0.0881	0.113	0.0494	0.0772
	(0.005)	(0.038)	(0.005)	(0.016)
Intercept	2.973	2.810	3.197	3.103
	(0.004)	(0.024)	(0.068)	(0.125)
N	656,904	3,677	656,904	656,904
R^2	0.004	0.007	0.124	0.582
Nationwide Sample	X		X	X
Chicagoland Sample		X		
Controls			X	X
Vehicle FE				X

Note: The table presents estimates the gender gap in log driving speed using data from the National Household Travel Survey. Dependent variable is average miles per hour driven on a single trip. Column 3 includes controls for household income bins, driver education bins, dummies for why the trip was taken, dummies for why the previous trip was taken, day of the week, hour of day, age dummies, MSA size bins, and whether the interstate was used on trip. Column 4 add individual vehicle fixed effects. Since each household only record trips on a single day, Column 4 only compares male and female speeds driven in the same vehicle on the same day. Standard errors in parentheses.